Decentralized Hidden Markov Modeling with Equal Exit Probabilities

Dongyan Sui
Fudan University
Shanghai, China
dysui22@m.fudan.edu.cn

Haitian Zheng
Fudan University
Shanghai, China
htzheng24@m.fudan.edu.cn

Siyang Leng
Fudan University
Shanghai, China
syleng@fudan.edu.cn

Stefan Vlaski
Imperial College London
London, United Kingdom
s.vlaski@imperial.ac.uk

***Abstract*—Social learning strategies enable agents to infer the underlying true state of nature in a distributed manner by receiving private environmental signals and exchanging beliefs with their neighbors. Previous studies have extensively focused on static environments, where the underlying true state remains unchanged over time. In this paper, we consider a dynamic setting where the true state evolves according to a Markov chain with equal exit probabilities. Based on this assumption, we present a social learning strategy for dynamic environments, termed Diffusion α-HMM. By leveraging a simplified parameterization, we derive a nonlinear dynamical system that governs the evolution of the log-belief ratio over time. This formulation further reveals the relationship between the linearized form of Diffusion α-HMM and Adaptive Social Learning, a well-established social learning strategy for dynamic environments. Furthermore, we analyze the convergence and fixed-point properties of a reference system, providing theoretical guarantees on the learning performance of the proposed algorithm in dynamic settings. Numerical experiments compare various distributed social learning strategies across different dynamic environments, demonstrating the impact of nonlinearity and parameterization on learning performance in a range of dynamic scenarios.**

***Index Terms*—adaptive learning, Bayesian inference, hidden Markov model, nonlinear dynamical systems, social learning**

I. Introduction and Related Work

Social learning refers to the process by which networked agents infer the underlying true state of the environment by gathering information and sharing beliefs. In multi-agent or social networks, Bayesian and non-Bayesian social learning models [1]–[5] have been extensively employed in economics, sociology, and engineering to characterize the behaviors of financial markets, social groups, and multi-agent systems [6]–[8]. Traditional models have mainly focused on static environments. Recent research has increasingly explored online social learning models in dynamic settings, such as adaptive social learning [9], [10] and diffusion hidden Markov modeling strategies [11], [12].

In this paper, we consider the online social learning problem, where a network of N agents labeled by $k = 1, \dots, N$ receive noisy observations/signals $\boldsymbol{\xi}_{k,i}$ (bold notation is used for random variables) of the evolving state at each time step $i \geq 1$. Their aim is to collectively estimate the underlying true state $\boldsymbol{\theta}_i^\star$ at each time instant i given the streaming observations $\boldsymbol{\xi}_{k,1}, \dots, \boldsymbol{\xi}_{k,i}$.

For simplicity, we assume that the true state $\boldsymbol{\theta}_i^\star$ belongs to a discrete set of M possible states $\Theta = \{\theta_0, \theta_1, \dots, \theta_{M-1}\}$. Each agent k assigns a belief to each state $\theta \in \Theta$ at each time step i, denoted by $\boldsymbol{\mu}_{k,i}(\theta)$. The belief characterizes the agents' confidence that θ is the underlying true state at time i and is a probability distribution over all possible states Θ, i.e., $\sum_{m=0}^{M-1} \boldsymbol{\mu}_{k,i}(\theta_m) = 1$, for all $i = 0, 1, \dots$ and $k = 1, \dots, N$. To avoid triviality, we assume that each agent's initial belief, $\boldsymbol{\mu}_{k,0}(\theta)$, is strictly positive for all $\theta \in \Theta$. Correct learning is said to occur at time i for agent k if the belief $\boldsymbol{\mu}_{k,i}(\theta)$ is maximized at the true state $\theta = \boldsymbol{\theta}_i^\star$. The observations $\boldsymbol{\xi}_{k,i}$ are independent random variables over time i conditioned on the true state $\boldsymbol{\theta}_i^\star$, taking values in the space Ξ_k. Given the underlying true state $\theta_i^\star$, the observations follow a probability density function $f(\cdot|\theta_i^\star)$, which implies that when the underlying state of the environment remains unchanged, the observations are independent and identically distributed (i.i.d.) random variables over time. Each agent k is equipped with a model that specifies the likelihood of the observations $\xi \in \Xi_k$ for each possible state $\theta \in \Theta$, denoted by $L_k(\xi|\theta)$.

The likelihood model $L_k(\xi|\theta)$, as a function of ξ, can be either a probability density function or a probability mass function, depending on whether ξ is continuous or discrete. To ensure the agents can successfully learn the underlying true state, we impose the following assumptions, which are also the typical assumptions in traditional social learning methods [2], [4], [5], [9], [13]:

Assumption 1 (Finiteness of KL Divergence)**.** For each pair of distinct states θ and θ', the Kullback–Leibler (KL) divergence [14] between $L_k(\xi \mid \theta)$ and $L_k(\xi \mid \theta')$ for any agent k satisfies $D_{\text{KL}}(L_k(\xi|\theta)||L_k(\xi|\theta')) < \infty$.

This assumption avoids trivial cases where a likelihood function model for a certain state completely dominates.

Assumption 2 (Global Identifiability of the Underlying True State)**.** The agents are collectively able to identify the true state uniquely:

$$\{\theta_i^\star\} = \Theta_i^\star = \bigcap_{k=1}^{N} \Theta_{k,i}^\star, \tag{1}$$

This work was conducted while D.S. was a visiting PhD student at Imperial College London with support from the China Scholarship Council (No. 202406100217). S.L. was supported by the National Natural Science Foundation of China (No. 12471457).

where $\Theta^\star_{k,i} = \arg\min_{\theta \in \Theta} D_{\text{KL}}(f(\cdot|\theta^\star_i)||L_k(\cdot|\theta))$.

The agents interact in a network. We denote by $A = [a_{\ell k}]$ the weight matrix of the network, which is assumed to be *nonnegative and left-stochastic*, i.e., $0 \le a_{\ell k} \le 1$, $\sum_{\ell=1}^N a_{\ell k} = 1$, $a_{\ell k} = 0$ for $\ell \notin \mathcal{N}_k$, where $\mathcal{N}_k$ denotes the neighborhood of agent k, with k itself being included.

Assumption 3 (Strong connectivity and aperiodicity)**.** The network of agents is strongly connected, and at least one node k in the network has a self-loop, i.e., $a_{kk} > 0$.

Under Assumption 3, the weight matrix A is a primitive matrix. According to the Perron-Frobenius theorem [15], there exists a Perron vector π satisfying:

$$A\pi = \pi, \quad \sum_{k=1}^N \pi_k = 1, \quad \pi_k > 0 \text{ for all } k = 1, \dots, N. \quad (2)$$

A. Diffusion α-HMM

We model the true state $\boldsymbol{\theta}^\star_i$ as a random variable, following a Markov chain. If the transition probabilities are denoted by $P = [p_{nm}]_{M\times M}$, where $p_{nm} = \mathbb{P}\left[\theta^\star_i = \theta_m \mid \theta^\star_{i-1} = \theta_n\right]$, then the optimal private belief update, based on the hidden Markov model (HMM), is given by:

$$\boldsymbol{\psi}_{k,i}(\theta_m) = \frac{\sum_{n=0}^{M-1} p_{nm}\boldsymbol{\mu}_{k,i-1}(\theta_n)L_k(\boldsymbol{\xi}_{k,i}|\theta_m)}{\sum_{\ell=0}^{M-1}\sum_{n=0}^{M-1} p_{n\ell}\boldsymbol{\mu}_{k,i-1}(\theta_n)L_k(\boldsymbol{\xi}_{k,i}|\theta_\ell)}. \quad (3)$$

This type of private belief update rule in social learning has been studied in [11], [12]. Observe that the optimal private belief update under a hidden Markov model involves the entries of the full state transition matrix P, which is frequently unknown and challenging to estimate in practice. In this work, we will instead study a simplified HMM-based update, which is derived under the assumption of equal exit probabilities for state transitions. Specifically, the true state transitions with a probability h, and when a transition occurs, the next state is chosen uniformly at random. Under these assumptions, the transition probability matrix simplifies to $p_{mn} = h/(M-1)$ if $m = n$, and $p_{mm} = 1 - h$.

A similar transition model was used in [10] to quantify the dynamics of Adaptive Social Learning in time-varying environments. Motivated by the update in (3), we adopt the following belief update rule:

$$\boldsymbol{\psi}_{k,i}(\theta_m) = \frac{((1-\alpha M)\boldsymbol{\mu}_{k,i-1}(\theta_m) + \alpha)L_k(\boldsymbol{\xi}_{k,i}|\theta_m)}{\sum_{n=0}^{M-1}((1-\alpha M)\boldsymbol{\mu}_{k,i-1}(\theta_n) + \alpha)L_k(\boldsymbol{\xi}_{k,i}|\theta_n)}, \quad (4)$$

where $\alpha = \frac{h}{M-1}$ represents the exit probability. This private belief update rule, referred to as the α-HMM, simplifies the inference problem to a single tunable hyperparameter α. Of course, for a general state transition matrix P, the simplified update rule (4) will be suboptimal compared to the optimal HMM filter (3). The advantage, on the other hand, is that this simplification enhances the theoretical tractability, and (4) relies only on a single parameter α, which quantifies the volatility of the underlying true state. Indeed, examining (4), we observe that α essentially controls the amount of weight placed on prior beliefs compared to the most recent observation. Note that when $\alpha = 0$, the above iteration degenerates into the classical Bayes' update:

$$\boldsymbol{\psi}_{k,i}(\theta_m) = \frac{\boldsymbol{\mu}_{k,i-1}(\theta_m)L_k(\boldsymbol{\xi}_{k,i}|\theta_m)}{\sum_{n=0}^{M-1}\boldsymbol{\mu}_{k,i-1}(\theta_n)L_k(\boldsymbol{\xi}_{k,i}|\theta_n)}. \quad (5)$$

On the other hand, when $\alpha = 1/M$, Eq. (4) reduces to:

$$\boldsymbol{\psi}_{k,i}(\theta_m) = \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_m)}{\sum_{n=0}^{M-1} L_k(\boldsymbol{\xi}_{k,i}|\theta_n)}. \quad (6)$$

In this case, each agent will negate prior beliefs and rely solely on the current private observation for learning. In this paper, since we focus on filtering in dynamic environments, we only consider the case where $0 < h < 1$ and consequently, $\alpha > 0$.

At time i, after the inference step (4), each agent k aggregates private beliefs from its neighboring nodes to form its current belief $\boldsymbol{\mu}_{k,i}$ using a geometrically weighted average:

$$\boldsymbol{\mu}_{k,i}(\theta_m) = \frac{\exp\left(\sum_{\ell\in\mathcal{N}_k} a_{\ell k}\log\boldsymbol{\psi}_{\ell,i}(\theta_m)\right)}{\sum_{n=0}^{M-1}\exp\left(\sum_{\ell\in\mathcal{N}_k} a_{\ell k}\log\boldsymbol{\psi}_{\ell,i}(\theta_n)\right)}. \quad (7)$$

Combining (4) and (7) together, we finally obtain the *Diffusion α-HMM strategy* for social learning in dynamical environment. Note that the strategy corresponds to a simplified form of the Diffusion HMM strategy from [11], [12]. In contrast to these works, we will not assume the Markov chain to be consistent with the dynamics driving $\boldsymbol{\theta}^\star_i$, and instead treat α as a tunable parameter akin to a step-size.

B. Dynamics of Log-belief Ratio in Steady State

For the purposes of analysis, we assume the environment remains in a single state over an extended period, resulting in fixed observation statistics. Without loss of generality, we assume that the underlying true state is $\theta_0 \in \Theta$, i.e., $\theta^\star_i = \theta_0$ for all $i \ge 1$. We aim to analyze the dynamics of the log-likelihood ratio of the belief on wrong and true states for any agent k, i.e.:

$$\boldsymbol{x}_{k,i}(\theta_m) \triangleq \log\frac{\boldsymbol{\mu}_{k,i}(\theta_m)}{\boldsymbol{\mu}_{k,i}(\theta_0)}, \quad m = 1, \dots, M-1. \quad (8)$$

It can be verified that the log-likelihood ratio $\boldsymbol{x}_{k,i}(\theta_m)$ in Diffusion α-HMM evolves as:

$$\boldsymbol{x}_{k,i}(\theta_m) = \sum_{\ell\in\mathcal{N}_k} a_{\ell k}\left(F_m(\boldsymbol{x}_{\ell,i-1}) + \log\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_m)}{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_0)}\right), \quad (9)$$

where

$$F_m(x_1, \dots, x_{M-1}) \triangleq \log\frac{(1-\alpha M)\exp(x_m) + \alpha + \alpha\sum_{n=1}^{M-1}\exp(x_n)}{1 - \alpha M + \alpha + \alpha\sum_{n=1}^{M-1}\exp(x_n)}, \quad (10)$$

$$\boldsymbol{x}_{k,i} = \left[\boldsymbol{x}_{k,i}(\theta_1), \dots, \boldsymbol{x}_{k,i}(\theta_{M-1})\right]^\top. \quad (11)$$

The following remark shows the connection between Diffusion α-HMM and the Adaptive Social Learning strategy from [9].

Remark 1. The nonlinear function (10) has the following properties:

$$F_m(0,\dots,0)=0, \tag{12}$$

$$\frac{\partial F_m}{\partial x_m}\Big|_{(0,\dots,0)}=1-\alpha M, \tag{13}$$

$$\frac{\partial F_m}{\partial x_n}\Big|_{(0,\dots,0)}=0,\quad \forall n\neq m. \tag{14}$$

By applying a multivariate Taylor expansion to $F_m(x)$ around $x=0_{M-1}$ up to the first-order term, for all $m=1,\dots,M-1$ we have $F_m(x_1,\dots,x_{M-1})=(1-\alpha M)x_m+o(\|x\|)$. Then, the linear approximation of the system (9) is given by:

$$\begin{aligned}\boldsymbol{x}_{k,i}(\theta_m)=&(1-\alpha M)\sum_{\ell\in\mathcal{N}_k}a_{\ell k}\boldsymbol{x}_{\ell,i-1}(\theta_m)\\&+\sum_{\ell\in\mathcal{N}_k}a_{\ell k}\log\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_m)}{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_0)}.\end{aligned} \tag{15}$$

The above equation resembles the evolution of the log-belief ratio in the Adaptive Social Learning (ASL) strategy in [9], which has the following form:

$$\begin{aligned}\boldsymbol{x}_{k,i}(\theta_m)=&(1-\delta)\sum_{\ell\in\mathcal{N}_k}a_{\ell k}\boldsymbol{x}_{\ell,i-1}(\theta_m)\\&+\delta\sum_{\ell\in\mathcal{N}_k}a_{\ell k}\log\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_m)}{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_0)},\end{aligned} \tag{16}$$

where δ is the step-size parameter. Equations (15) and (16) both apply a discount to the information from the previous time step in the private belief update step. It is worth noting that another variant of ASL [13], [16] exhibits the same log-belief ratio dynamics as in (15), with $\delta=\alpha M$. Further comparison of the performance among the Diffusion α-HMM, linearized Diffusion α-HMM, and ASL will be illustrated in numerical experiments.

It can be seen that the recursion (9) is both *nonlinear* and *stochastic*. To facilitate analysis, we introduce the following deterministic dynamical reference system, where stochastic quantities are replaced by their expected values:

$$\hat{x}_{k,i}(\theta_m)=\sum_{\ell\in\mathcal{N}_k}a_{\ell k}\left(F_m(\hat{x}_{\ell,i-1})-d_\ell(\theta_m)\right), \tag{17}$$

where

$$\begin{aligned}d_k(\theta_m)\triangleq&D_{\text{KL}}(f(\cdot|\theta_0)||L_k(\cdot|\theta_m))-D_{\text{KL}}(f(\cdot|\theta_0)||L_k(\cdot|\theta_0))\\=&-\mathbb{E}\left[\log\frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_m)}{L_k(\boldsymbol{\xi}_{k,i}|\theta_0)}\right].\end{aligned} \tag{18}$$

Assumption 2 ensures that $d_k(\theta_m)\geq 0$ and for all $m=1,\dots,M-1$ there exists at least one $k=1,\dots,N$ such that $d_k(\theta_m)>0$. The quantity $d_k(\theta_m)$ quantifies agent k's ability to distinguish between an incorrect state θ_m and the true state θ_0. We refer to this measure as the identifiability of agent k with respect to θ_m.

In this paper, we first analyze the fixed point of the dynamical reference system (17) under steady-state conditions. We then establish the convergence of the reference system to its fixed point. Finally, under additional assumptions on noise and identifiability, we derive an estimate for the error probability.

II. Convergence Analysis

For a discrete dynamical system $x_i=T(x_{i-1})$, a fixed point corresponds to an equilibrium state of the system, such that $T(x^\infty)=x^\infty$. We begin by characterizing x_k^∞ for the reference system (17).

Lemma 1. The fixed points of the dynamical reference system (17) exist, and satisfy $\hat{x}_k^\infty(\theta_m)<-\sum_{\ell\in\mathcal{N}_k}a_{\ell k}d_\ell(\theta_m)$.

Proof. Proof omitted due to space limitations. □

The following theorem shows the convergence of (17) to its unique fixed point provided that the underlying state is constant.

Theorem 1. When $0<\alpha<1/M$ and the underlying true state $\boldsymbol{\theta}_i^\star$ remains constant, $\hat{x}_{k,i}(\theta_m)$ defined in (17) will converge to a unique fixed point, i.e.,

$$\lim_{i\to\infty}\hat{x}_{k,i}(\theta_m)=\hat{x}_k^\infty(\theta_m),\quad\forall 1\leq k\leq N,\ 1\leq m\leq M-1. \tag{19}$$

Proof. Proof omitted due to space limitations. □

Theorem 1 establishes that the dynamical reference system converges to a unique fixed point, whose value is upper-bounded by a weighted average of the identifiability of neighboring agents, referred to as the neighborhood identifiability. Thus, neighborhood identifiability guarantees the learning capability of agent k under steady-state conditions.

In the following, for the purpose of further theoretical analysis, we make the following assumption, which requires the log-likelihood to be bounded almost surely.

Assumption 4 (Bounded Log-Likelihood Ratio)**.** There exists a positive constant C such that:

$$\max_{k=1,\dots,N}\max_{m=1,\dots,M-1}\sup_{\boldsymbol{\xi}\in\Xi_k}\left|\log\frac{L_k(\boldsymbol{\xi}|\theta_m)}{L_k(\boldsymbol{\xi}|\theta_0)}+d_k(\theta_m)\right|\leq C,\ a.s.. \tag{20}$$

To evaluate the learning performance of Diffusion α-HMM, we introduce an important metric: the *instantaneous error probability*, defined as:

$$p_i^e\triangleq\mathbb{P}\left[\,\exists\, k=1,\dots,N \text{ and } \theta_m\neq\theta_0,\ \text{s.t. } \boldsymbol{\mu}_{k,i}(\theta_m)\geq\boldsymbol{\mu}_{k,i}(\theta_0)\right]. \tag{21}$$

From this definition, it is clear that the instantaneous error probability quantifies the probability that any agent fails to correctly identify the underlying true state at a given time i during the online social learning process.

Due to the nonlinearity and stochasticity of the original system, directly computing the error probability is challenging. Here, we provide an analytical result under the assumption of sufficient neighborhood identifiability.

Theorem 2. When $\alpha < 1/M$ and the underlying true state $\boldsymbol{\theta}_i^\star$ remains constant, if for all $k = 1, \dots, N$ and $m = 1, \dots, M-1$,

$$\sum_{\ell \in \mathcal{N}_k} a_{\ell k} d_\ell(\theta_m) > C, \tag{22}$$

then the instantaneous error probability for the diffusion $\alpha-$HMM algorithm satisfies:

$$\limsup_{i\to\infty} p_i^e \leq \frac{C}{-\alpha \overline{x}^\infty}, \tag{23}$$

where $\overline{x}^\infty = \max_{k=1,\dots,N} \max_{m=1,\dots,M-1} \hat{x}_k^\infty(\theta_m) < 0$, C is as defined in (20).

Proof. Proof omitted due to space limitations. □

Theorem 2 establishes that, under certain conditions on neighborhood identifiability, the upper bound of the instantaneous error probability converges to a fixed value. This result implies that within the Diffusion α-HMM framework, the probability of erroneous learning cannot be guaranteed to asymptotically approach zero, even in steady-state conditions, which contrasts with the behavior predicted by Bayes' formula [5]. However, by sacrificing some learning accuracy, the Diffusion α-HMM framework significantly enhances adaptability, as further demonstrated in the numerical experiments.

The steady-state error probability can be mitigated in one of two ways: (1) Shifting the fixed point $\hat{x}_k^\infty(\theta_m)$ of the deterministic system further away from zero. As shown in Lemma 1, improving agents' identifiability $d_k(\theta_m)$ facilitates this shift. (2) Reducing observation noise, thereby increasing the informativeness of the received signals.

III. Numerical Experiments

In the numerical experiments, we consider a network of $N = 5$ agents attempting to infer the evolving state from a set of $M = 3$ possible states. Different network topologies are examined, including a fully connected network and a strongly connected network. The weight matrix is randomly initialized while ensuring it satisfies the conditions of a primitive matrix. The true distribution of the observation $\boldsymbol{\xi}_{k,i}$, $f(\cdot|\theta_i^\star)$, follows a normal distribution $\mathcal{N}(\theta_i^\star, \sigma^2)$, where $\theta_i^\star \in \Theta = \{0, 1, 2\}$. Each agent k employs a likelihood model $L_k(\cdot|\theta)$, which is also modeled as a normal probability density function with a standard deviation of σ and a mean specified in Table I. As observed from Table I, some agents are unable to independently distinguish among the three states. However, for any pair of states, there exists at least one agent capable of differentiation, which satisfies Assumption 2.

We compare the performance of three algorithms: the Diffusion α-HMM introduced in this paper (Eq. (4) and Eq. (7)), the linearized Diffusion α-HMM (15), and adaptive social learning [8]. Specifically, consider the private belief update given by:

$$\boldsymbol{\psi}_{k,i}(\theta_m) = \frac{\boldsymbol{\mu}_{k,i-1}^{1-\delta_1}(\theta_m) L_k^{\delta_2}(\boldsymbol{\xi}_{k,i}|\theta_m)}{\sum_{n=0}^{M-1} \boldsymbol{\mu}_{k,i-1}^{1-\delta_1}(\theta_n) L_k^{\delta_2}(\boldsymbol{\xi}_{k,i}|\theta_n)}. \tag{24}$$

When $\delta_1 = \delta_2 = \delta$, this update rule corresponds to the adaptive social learning (ASL) strategy in [9]. When $\delta_1 = \alpha M$ and $\delta_2 = 0$, it represents the linearized Diffusion α-HMM.

In the first scenario, the underlying true state evolves according to a Markov chain with equal exit probabilities, which is the same mechanism on which our proposed algorithm is based. The true exit probability is set to $\alpha_0 = 0.1$, and we tune the parameter α with different fixed values of σ in the Diffusion α-HMM, linearized Diffusion α-HMM, and ASL (where δ is set as αM). Due to the ergodicity of the underlying stochastic process governing the state evolution, we compute the probability of successfully tracking the true state over 50,000 time steps. The comparison results are presented in Fig. 1.

A horizontal comparison in Fig. 1 reveals that although the learning accuracy is generally higher in the fully connected network, in the strongly connected network, even agent 5, which lacks distinguishing ability, can still correctly infer the true state most of the time due to information aggregation from other agents. A vertical comparison shows that as the standard deviation of noise σ increases, the learning accuracy decreases. When comparing the three algorithms, we observe that under low noise conditions, the highest learning accuracy follows the order: Diffusion α-HMM $>$ Linearized Diffusion α-HMM $>$ Adaptive Social Learning. Moreover, the first two methods exhibit superior performance over a broader range of parameter values compared to ASL. Additionally, Diffusion α-HMM consistently outperforms the linearized model, demonstrating the benefits of nonlinearity in improving learning accuracy. However, when the noise level is high, ASL achieves slightly higher maximum learning accuracy than both Diffusion α-HMM and Linearized Diffusion α-HMM in the fully connected network, highlighting the role of step-size scaling in mitigating noise effects.

In the second scenario, we assume that every T_0 iterations, the underlying true state is randomly selected from Θ. Notably, under this assumption, the state evolution does not follow a Markov chain. In the strongly connected network, we evaluate the performance of the three algorithms under different transition intervals T_0 and noise standard deviations σ. From Fig. 2, we observe that despite the change in the state evolution mechanism, the relative performance of the three algorithms remains similar to that in Scenario 1. This further highlights the role of the nonlinear inference step in dynamic environments. Additionally, when the environment

TABLE I
LIKELIHOOD MODEL CONFIGURATION FOR THE AGENTS.

Agent k	Likelihood model $L_k(\cdot\|\theta)$		
	$\theta_0 = 0$	$\theta_1 = 1$	$\theta_2 = 2$
1	0	1	2
2	0	1	1
3	0	0	2
4	0	1	0
5	0	0	0

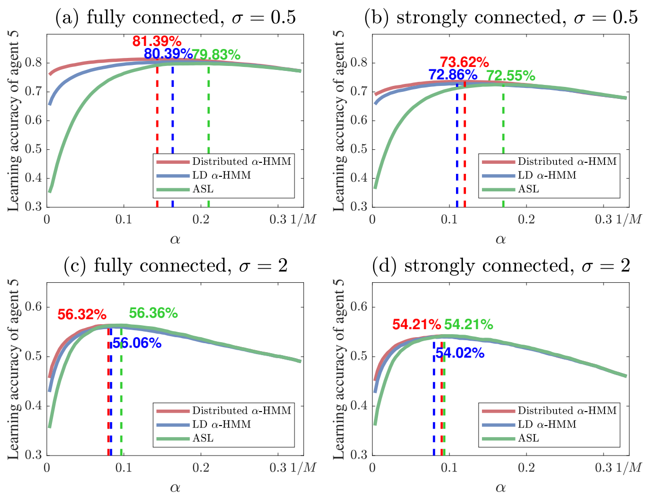

Fig. 1. Comparison of the learning accuracy of agent 5, which has no distinguishing capability, under different network topologies and algorithms in Scenario 1. The accuracy is evaluated as a function of α for different fixed values of σ. LD α-HMM refers to the linearized Diffusion α-HMM.

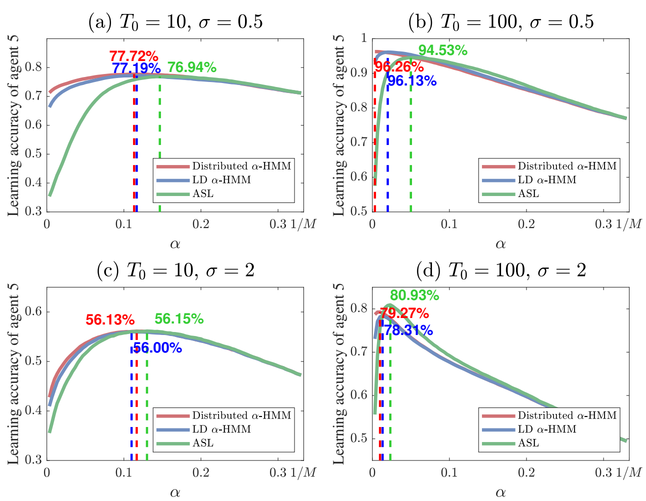

Fig. 2. Comparison of the learning accuracy of agent 5, which has no distinguishing capability, under different transition intervals T_0 and algorithms in Scenario 2. The accuracy is evaluated as a function of α for different fixed values of σ. LD α-HMM refers to the linearized Diffusion α-HMM.

is less volatile (e.g., Fig. 2(b)), the performance of ASL for small δ deteriorates. This indicates that a small δ applied to observational data may hinder the learning capability of the multi-agent system.

Fig. 3 illustrates the evolution of agent 5's belief on the true state under Scenario 2, comparing the three algorithms with their respective optimal parameter values, $\alpha^\star$, in two different settings. From the figure, it can be observed that when the observation noise σ is small and the environment is more dynamic (i.e., smaller T_0), the Diffusion α-HMM exhibits a faster adaptation rate. Conversely, when the observation noise is large and the environment is more stable, ASL can mitigate the impact of observation noise by selecting a smaller α (corresponding to a smaller step-size δ), resulting in a

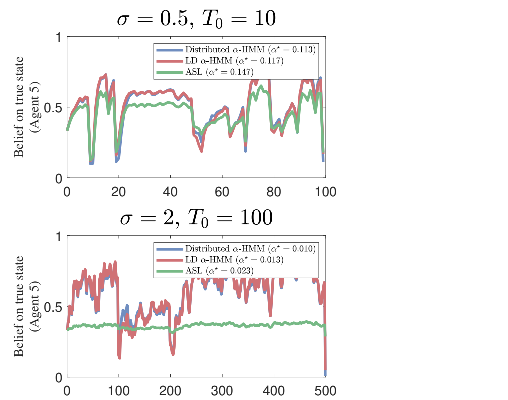

Fig. 3. Comparison of agent 5's belief evolution on the true state across the three algorithms under two different settings.

smoother belief evolution. However, in such cases, the beliefs across different states tend to be closer to each other.

REFERENCES

[1] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[2] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-bayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.

[3] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "A theory of non-bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, 2018.

[4] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.

[5] A. Lalitha, T. Javidi, and A. D. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp. 6161–6179, 2018.

[6] C. Chamley, *Rational herds: Economic models of social learning*. Cambridge University Press, 2004.

[7] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, pp. 3–49, 2011.

[8] V. Bordignon, S. Vlaski, V. Matta, and A. H. Sayed, "Learning from heterogeneous data based on social interactions over graphs," *IEEE Transactions on Information Theory*, vol. 69, no. 5, pp. 3347–3371, 2022.

[9] V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning," *IEEE Transactions on Information Theory*, vol. 67, no. 9, pp. 6053–6081, 2021.

[10] M. Khammassi, V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning for tracking rare transition markov chains," in *2024 32nd European Signal Processing Conference (EUSIPCO)*, 2024, pp. 1032–1036.

[11] M. Kayaalp, V. Bordignon, S. Vlaski, and A. H. Sayed, "Hidden markov modeling over graphs," in *2022 IEEE Data Science and Learning Workshop (DSLW)*. IEEE, 2022, pp. 1–6.

[12] M. Kayaalp, V. Bordignon, S. Vlaski, V. Matta, and A. H. Sayed, "Distributed bayesian learning of dynamic states," *arXiv preprint arXiv:2212.02565*, 2022.

[13] V. Bordignon, V. Matta, and A. H. Sayed, "Socially intelligent networks: A framework for decision making over graphs," *IEEE Signal Processing Magazine*, vol. 41, no. 4, pp. 20–39, 2024.

[14] M. Thomas and A. T. Joy, *Elements of information theory*. Wiley-Interscience, 2006.

[15] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge university press, 2012.

[16] V. Matta, V. Bordignon, A. H. Sayed *et al.*, "Social learning: Opinion formation and decision-making over graphs," in *EURASIP–Now Publishers Open Access Book Series on Information and Learning Sciences*. Now Publishers, Inc., 2025.